Title: HPCC Mirroring and IP Load Balancing Project

Purpose:

The goal of this project is to demonstrate two technical capabilities:

1. The ability to replicate the content of one server (source) to other geographically distributed servers (mirrors) and to maintain the consistency between the mirrored servers and the source server.

2. Provide a capability to dynamically and intelligently balance the IP request between the source server and the mirror servers in order to optimize the "load" on any single server and provide robust accessability to the information that these servers provide.

Overview of the Project:

This is a cooperative project, sponsored by the NOAA High Performance Computer and Communications (HPCC) office, involving the Charleston National Weather Service(NWS), NOAA Network Information Center (NIC) and the Network Operations Centers (NOCs) in Boulder, CO. and Silver Spring, MD..

This particular project involved two major components:

I. The mirroring of the NWS Charleston S.C. web site (wchs.csc.noaa.gov) at two locations:

- The Network Operations Center (NOC) in Boulder, CO.; and
- The NOAA Network Information Center (NIC) in Suitland, MD..

II. Implementation of a capability to dynamically and intelligently balance the IP request between the source server and the mirror servers.

Diagram:

<Insert Diagram here>

I.) Content Mirroring and Updating

The mirroring and updating was accomplished using two freeware UNIX software programs RSYNC and SSH. RSYNC is a file transfer program for Unix systems. RSYNC uses the "rsync algorithm" which provides a very fast method for bringing remote files into sync. It does this by sending just the differences in the files across the link, without requiring that both sets of files are present at one of the ends of the link beforehand. SSH(Secure Shell) is a program to log into another computer over a network, to execute commands in a remote machine, and to move files from one machine to another. It provides strong authentication and secure communications over insecure channels.

A.) Instructions for Setting Up SSH and RSYNC:

1. Original Site Requirements:

Install ssh v1.x
Add RSA host public key identification of distributed mirrors to ~/.ssh/known_hosts file
Create RSA key via ssh-keygen with null passphrase
Install rsync v2.x
Create mirror user to transfer files (must have read access of files to be mirrored)
Requires mirrored content to be in a tree structure
Standard CGI scripts
Add rsync script to crontab at desired frequency to push content out to the distributed mirrors
Set site to DNS name of original site
Set lpath to root of directory to be mirrored
Set lrsync to location of local rsync
Set lssh to location of local ssh

2. Distributed Mirror Requirements:

Install ssh v1.x
Add RSA public key from original site's identity_pub to ~/.ssh/authorized_hosts file
Install rsync v2.x in /usr/local/bin (we will not be using rsync as a server so there's no need for a rsync.conf file)
Create mirror user (mirror) to transfer files.
Create mirror directory /web/mirror/$site where $site is the DNS name of the site to be mirrored

Install a Virtual Host capable web server (i.e. Apache)
For websites, add virtual hosting configuration to HTTPD server configuration file, i.e.

```
<VirtualHost 140.172.XXX.XX>
  ServerAdmin webmaster@wchs.test.noaa.gov
  DocumentRoot /web/mirror/wchs.test.noaa.gov
  ServerName wchs.test.noaa.gov
  ErrorDocument 403 /error403.html
  ErrorDocument 404 /error404.html
  ErrorLog var/log/wchs.test.noaa.gov/error_log
  TransferLog var/log/wchs.test.noaa.gov/access_log
</VirtualHost>
```

Install a Virtual Host capable FTP server (i.e. wu-ftpd)

For FTP sites, add virtual hosting configuration to the configuration files, i.e.

```
# Virtual Server at 140.172.XXX.XX
virtual 140.172.XXX.XX root    /web/mirror/ftp.test.noaa.gov/ftp
virtual 140.172.XXX.XX banner  /web/mirror/ftp.testnoaa.gov/banner.msg
virtual 140.172.XXX.XX logfile /web/mirror/ftp.test.noaa.gov/xferlog
```

II.) IP Load Balancing

This project used CISCO's Distributed Director to provide the dynamic and intelligent load distribution of IP requests between the source and mirrored servers. The Cisco Distributed Director provides the ability to perform load distribution in a sophisticated manner that takes server availability, relative client-to-server topological proximities ("distances"), and client-to-server link latency into account to determine the "best" server. This means that users need only a single subdomain name or Universal Resource Locator (URL)-embedded hostname for accessing a distributed set of servers. This eliminates the need for end-users to choose a server from a list of possible sites. The Cisco Distributed Director leverages the intelligence in the network to automatically, dynamically, and efficiently pick the "best" server for the user, using a single hostname or DNS subdomain name.

The Distributed Director is configured to act as the primary DNS name server for a specific subdomain.

A.) Distributor Director Sample Configuration:

```
;Add new mirrored host name and IP addresses
 ip host wchs.test.noaa.gov 140.90.XXX.XX 140.172.XXX.XX
;Add SOA record for new mirror
 ip dns primary wchs.test.noaa.gov soa dd.boulder.noaa.gov dns.boulder.noaa.gov
21600 900 60 0
;Add server keepalive check
 ip director hosts wchs.test.noaa.gov connect 80 interval 900
;Add mirror to access-list for the return of DNS responses
 ip director accesss-list 1 permit wchs.test.noaa.gov
; Assign weights and priorities to metrics
 ip director hosts wchs.test.noaa.gov weights ran 100
 ip director hosts wchs.test.noaa.gov priority ran 1
            ......
```

B.) Primary Name Server Sample Configuration:

Delegate domain to Distributed Directors, i.e. in the test.noaa.gov zone file

```
wchs    IN    NS    dd.suitland.noaa.gov
```

IN NS dd.boulder.noaa.gov

C.) Router Configuration:

IOS - 11.3(2)T or later must be running

Turn on DRP agent:

```
ip drp server
```

Enable security

```
access-list 1 permit 140.90.231.30 140.172.6.240
                     (NIC DD)      (Boulder DD)
access-list 1 deny any
```

Ensure router accepts DRP queries

```
ip drp access-group 1
```

Set up a key chain

```
key chain <some name>
key 1
key-string <some string>
exit
```

Enable DRP authentication key chain

```
ip drp authentication key-chain   <some name>
```

Verify key chain configuration

```
Show ip drp
```

III.) Test Plan

The testing of the Distributed Directors employed a phased approach.

Phase I:

The wchs.csc.noaa.gov web site was mirrored, using RSYNC and SSH, at 2 sites, one at the NOC in Boulder, CO. and the other at the NIC in Suitland, MD.. A test domain name was then created, wchs.test.noaa.gov, for the Boulder and NIC mirror sites.

2 Distributed Directors were used, one at the NOC in Boulder, CO. and the other at the NIC in Suitland, MD..

Both Distributed Directors were configured to use a single metric (Random) to distribute the HTTP requests randomly between the two mirrored servers.

A group of 12 users was asked to participate.

The test lasted approximately 8 hours.

The test involved users opening up a window on their browsers and accessing the Wchs.test.noaa.gov web site. Users were told to leave their browsers on all day. The web page was set to automatically refresh every minute.

Phase II:
This test involved the use of the CISCO's Director Response Protocol (DRP) agent software. Two CISCO routers, one at the Boulder NOC and the other at the NIC were set up to use the DRP agent software.

A.) Distributed Director Metrics:

The two Distributed Directors were set up to use the DRP Round Trip Time (drp-rtt) metric as the primary determinant to distribute traffic between the Boulder and NIC web servers. The random metric was also used as a secondary determinant.

B.) Additional Distributor Director Sample Configurations:

```
ip director server 140.90.XXX.XX drp-association 207.24.XXX.X
ip director server 140.172.XXX.XX drp-association 140.172.X.XXX
ip director hosts wchs.test.noaa.gov weights ran 10
ip director hosts wchs.test.noaa.gov weights drp-r 90
ip director hosts wchs.test.noaa.gov priority ran 2
ip director hosts wchs.test.noaa.gov priority drp-r 1
```

The test domain, wchs.test.noaa.gov, was again used in this test.

A group of 12 users was asked to participate.

The test lasted approximately 8 hours.

The test involved users opening up a window on their browsers and accessing the wchs.test.noaa.gov web site. Users were told to leave their browsers on all day. The web page was set to automatically refresh every minute.

During the test one of the web servers was intentionally disconnected to see if the Distributed Directors would adjust the flow of traffic to the remaining server.

Phase III:

This test will involve:

1. Mirror a NESDIS FTP site
2. Use DDs to load balance FTP traffic
3. Perform a test of the DD to see how it balances both HTTP And FTP traffic together.

IV.) Test Results

Phase I Test:

Results were measured at two levels. Web server statistics were captured on both the NIC and Boulder Web servers. Both Distributed Directors were checked to see how they distributed the HTTP requests they received.

Web Server statistics:

See this URL to see a breakdown of the combined web statistics from both the NIC and Boulder web servers.

Approximately 74% of the HTTP requests were recorded by the NIC web server and approximately 26% were recorded by the Boulder web server.

11 unique hosts were identified in the web statistics. 45% or 5 hosts, at some point in time during the test, were directed to both web servers. 55% or 6 hosts were always directed to the same web server.

Distributed Director Statistics:

The Distributed Director at the NIC directed 293 requests during the test period. 124 or 42% of the requests were directed to the Boulder web server. 169 or 58% of the requests were directed to the NIC web server.

The Distributed Director at Boulder directed 39 requests during the test period.

14 or 36% were directed to the NIC web server and 25 or 64% were directed to the Boulder web server.

Phase II Test:

Configuration Changes:

Two routers were configured to use the CISCO Director Response Protocol (DRP) agent software. One at the NOC in Boulder, CO. and the other at the NIC in Suitland, MD.. Both Distributed Directors were configured to use two metrics to distribute the HTTP requests between the two mirrored servers. The two metrics used were :
1. DRP-RTT - A measure of the round trip time between the client's DNS server and each mirror server. The DRP-RTT metric was given a 90% weighting factor and a priority of 1.
2. Random - The same random metric used in the phase one test. The random metric was given a 10% weighting factor and a priority of 2. If the RTT metric between each server were the same then each server was considered equal and one would be randomly chosen to answer the request.

The NIC's mirror web server was intentionally taken off line for a 15 minute period to see how the Distributed Director would handle this change.

Length of test: approximately 8 hrs.

Test Results:

Results were measured at two levels. Web server statistics were captured on both the NIC and Boulder Web servers. Both Distributed Directors were checked to see how they distributed the HTTP requests they received.

Web Server statistics:

The web statistics from the NIC and Boulder web servers.
Approximately 58% of the HTTP requests were handled by the NIC web server and approximately 42% were handled by the Boulder web server.
22 unique hosts were identified in the web statistics. 40% of the HTTP request, at some point in time during the test, were directed to both web servers. 60% of the HTTP request were always directed to the same web server.

Distributed Director Statistics:

The Distributed Director at the NIC directed 129 requests during the test period.
61 or 47% of the requests were directed to the Boulder web server.
68 or 53% of the requests were directed to the NIC web server.

The NIC monitored 29 request handled by the our Distributed Director, from 12:30 - 4:00 P.M..25 requests were indeed routed to the web server with the lowest measured round trip time.2 requests were routed using the random metric since the round trip times of the two servers were considered equal.

1 request was answered from the Distributed Director's Cache (not sure which server it chose).
4 periodic DRP error messages were noted by the NIC's Distributed Director indicating a problem with the DRP agent software on the NIC router. These errors did not seem to have any effect on any of the traffic distribution. CISCO has been made aware of these errors and is currently looking into it. Note the IOS used on the router is not a General Deployment (GD) version.
When the NIC web server was taken off line for 15 minutes, the Distributed Director was able to detect that the NIC server was no longer available. 5 request came in to the Distributed Director during this time and all 5 were sent to the Boulder web server. When the NIC server was brought back on line the Distributed Director detected it and began to route request to the NIC server again.

The NIC Distributor Director received a query for ACCSTAFF.ACC.NOS.NOAA.GOV. Not sure why this request was sent to the director.

The number of hits recorded by the webservers was not equal to the number of requests that were recorded by the Distributed Director.

Phase III Test:

Late November or early December time frame.

V.) Conclusions (To Date):

The initial test was successful in demonstrating that the Distributed Director is capable of directing IP requests over a wide area network.

Factors that may have influenced the initial test results:

- Most of the requests came from the 140.90 network which is where the NIC's web server and Distributed Director are located.

- The caching period of individual workstations may have had some effect

The second test, using both the Round Trip Time and Random metrics, appeared to distribute the load between the two mirror servers to almost a ratio of 50/50 (58/42). The traffic distribution measured on the NIC's distributed Director also indicated an almost 50/50 (68/61) ratio of request distribution between the two mirror servers. Using the random metric alone in the phase one test resulted in almost a 70/30 (74/26) distribution. So one might conclude that a more even distribution of sever load results from using the Distributed Director employing the Round Trip Time metric.

60% of the web traffic was directed to the same server vs. 31% from the phase one test. This change probably resulted because traffic was routed to the server with the lowest RTT and the RTT values probably did not vary much throughout the day. Other factors that came into play were occasions when the RTT values were equal, the NIC server being taken off line and the caching of DNS information by each workstation/PC and downstream name server.

It appears that the discrepancy between the number of hits recorded by the Distributed Director and the webservers is due to caching done by workstations and downstream name servers. This caching reduces the effectiveness of the Distributed Director and needs to be investigated further. Early versions of BIND (<4.9) do not handle DNS resource records with TTLs of 0 consistently.

A factor that will come into play should we make the Distributor Directors operational is that the version of the router IOS that is required is not currently a GD release. Most organizations do not wish to run a non-GD release IOS on their operational routers.

The Distributed Director appears to consistently route request to the server with the lower Round Trip Time measure.

The Distributed Director is able to dynamically detect when a server is no longer available and to route request to the remaining active server without any operator intervention. It is also able to dynamically detect when a server is again available and routes traffic to that server without any operator intervention.

From the results measured so far in the Phase I and II tests it does appear that the Distributed Director provides some value added capabilities above what is available from using standard Round Robin DNS.